

Sergei Ponomarev

Demand the AI Receipt

How consumers can get value, not headaches, from corporate AI

Five demands for corporate AI:

1. Check whether AI actually makes your life better.
2. Demand the AI receipt.
3. Demand the AI service passport.
4. Demand the AI policy.
5. Change the system to fit you.

aibusiness.vc

2026

Contents

Introduction. The Rules Are Being Written Right Now	3
Chapter 1. Not All AI Is Equally Useful.....	5
1.1. Check Three Things.....	5
1.2. Consumer Needs. Demand Better.....	7
1.3. Demand Plain, Human Language	9
Chapter 2. Demand the AI Receipt.....	12
2.1. The Receipt as the Gateway to Justice	12
2.2. What an AI Receipt Should Contain	14
Chapter 3. Demand the AI Service Passport.....	17
3.1. The Passport as a Yardstick of Quality.....	17
3.2. What the Passport Should Contain	18
Chapter 4. Demand the AI Policy.....	21
4.1. Why a Company Needs an AI Policy	21
4.2. What the AI Policy Should Contain.....	22
Chapter 5. Change the System to Fit You.....	26
5.1. First, Stand Up for Yourself.....	26
5.2. From Feedback to Co-authorship	27
Conclusion. The Rules Depend on Us	30
About the Author	32

Introduction. The Rules Are Being Written Right Now

Right now companies are deciding how artificial intelligence will serve us, consumers, for all the years ahead.

How AI bots will talk to us. How many times we will have to repeat the same thing. How frustrated we must become before the bot calls in a human, and whether it will call one at all. What personal data companies will collect and how they will process it. How AI makes decisions that affect us. Whether we will be able to appeal those decisions.

Most of these rules are being written right now, on the basis of millions of daily interactions between people and corporate AI agents.

Companies pick models, hire developers, count the savings, lay off staff. All of that is understandable. But amid these calculations the main question often gets lost:

And why do we, consumers, need your AI?

What does it actually give us? We are promised that artificial intelligence will make service faster, simpler and cheaper. Wonderful. Let us check.

Has the number of steps gone down? Has the time shrunk? Has it become simpler to get a result? Have our costs dropped? Are we not now spending forty minutes going in circles with a chatbot instead of having a five-minute conversation with a helpful human agent?

An instant reply does not yet mean fast service. A bot can answer in a second and still mess with your head for half an hour. Round-the-clock service is not worth much if the service is bad.

If, after AI is rolled out, we spend more time, take more steps, suffer more stress and put in more effort, that is not innovation for the consumer. It is the company shifting its costs onto our shoulders. And no, we are not obliged to learn how to talk properly to every new bot. If a bot does not understand me, that is the system's problem, not mine.

We need to talk about this. Write about it. Demand it.

This brochure is not aimed against corporate artificial intelligence. On the contrary, I want to see more good AI: it can save time, help us make sense of complicated services, work at night and on weekends, and take on dreary routine. But useful AI will not emerge through companies' efforts alone. Consumers must make their voices heard and state their demands and preferences.

Why I Am Writing About This

I am not a lawyer with a ready answer for every country, industry and situation. My professional lens is a little different: the quality, openness and convenience of services for consumers.

For more than twenty years I have studied how digital and government services work in real life. For seven years I ran monitoring programmes and independent service-quality assessments, conducted mystery shopping, developed information-transparency standards, worked with government and human-rights organizations, taught, and advised entrepreneurs on service openness and quality.

In recent years I have been working on artificial intelligence, building agents myself and creating useful services. But in this brochure I am first of all the same consumer as you. Just a little more immersed in the subject and used to asking uncomfortable questions.

I suggest we ask them together. All the more so because the law is on our side. In Europe, from 2 August 2026, the transparency requirements of [Article 50 of the EU AI Act](#) start to apply. In particular, providers of systems intended for direct interaction with people must make sure a person is informed that they are talking to an AI, where this is not obvious. For certain kinds of systems and content, additional obligations apply.¹

Not everything I ask for in this brochure is already written into law as a company obligation. And that changes nothing. I am setting a consumer frame: what we have the right to want from corporate AI. How it maps onto the law depends on the country, the industry and the specific service, and let the companies' lawyers worry about that. We demand not because we found the right article, but because we pay. If service is inconvenient, unclear and disrespectful, we do not need the law's permission to say so out loud and walk away to someone who does it better.

This brochure is a consumer framework, not legal advice. Applicable rights and duties vary by jurisdiction, sector, the data involved, the company's role, and the use case. Legal information was checked on 22 July 2026.

We vote with money, time, attention, data and loyalty. There are not one or ten of us. There are billions of us. We choose companies, recommend them to other people, leave reviews, file complaints and walk away to competitors.

Where the law already gives us a right, we can demand it. Where the right is not yet formulated, we can build consumer practice.

This brochure is built around five simple actions.

First, work out whether AI makes our lives better.

Then learn to demand the AI receipt.

Demand that AI services meet clear standards.

Demand that the company disclose its AI policy.

Learn to secure changes and take part in creating the new rules.

Every chapter will end with a ready-made text you can send. You can send it to a company almost right away, inserting the name of the service and the details of your own case.

I will help you protect your personal interests: your time, money, data and peace of mind.

Right now there is a rare window of opportunity for this. Corporate AI is only just taking its place behind the counter. The rules have not yet settled. Companies are searching for answers themselves and often do not know how to build the right relationship with us, the consumers.

So let us not wait until everything is already decided. Let us take part now.

Chapter 1. Not All AI Is Equally Useful

Companies love to tell us that artificial intelligence makes service faster, simpler, and more convenient. We will not argue. We will check.

Not by a slick presentation, not by how fast the bot replies, and not by how much the company saves. We will check what changed for us personally: how many actions we take, how much time we spend, and how much it all costs us.

If these numbers went down, the AI is genuinely useful. If they did not change, the company simply automated its own work. If they went up, then under the guise of innovation someone else's work and costs have been shifted onto us.

In this chapter we will learn to tell useful AI from harmful AI. It is fairly simple.

1.1. Check Three Things

The clearest consumer assessment comes from contrast. Say for years you received a service with the help of a staff member (you shopped, made returns, and so on), and then the company rolled out an AI bot. Recall how it used to work, then walk the same path with the new AI bot.

Compare three parameters.

The first number: your actions.

Count everything you had to do from the moment of "I need this" to the moment of "my question is resolved."

Find the right section on the site: step 1. Register: step 2. Confirm your email. Choose a request topic. Explain the problem to the bot. Rephrase the question. Upload a document. Upload the document again, because the first one vanished somewhere. Ask the bot to connect you to an agent. Tell the agent everything from the start, because the bot did not pass on the chat.

Every tap, transition, repeat, and switch is your action, or step.

Do not count only what happens inside the chat. A company may proudly report that its bot resolves a question in three messages. But before those three messages you spent ten minutes hunting for the chat itself, registering, and proving you are not a robot. And after talking to the bot you still called a staff member.

You need to count the whole journey: from the moment a need arises until the issue is resolved.

Suppose you used to call a staff member, give your order number, and change the delivery date. Three actions.

Now you log in to your account, pass two-factor authentication, look for the chat, choose a city, choose a topic, enter the order number, explain to the bot what you want, wait for confirmation, then move to an agent, tell them everything again, and get a confirmation. Eleven actions.

The second number: your time.

Time it from the moment you started solving the problem to the moment it is actually solved.

Not to the first reply. Not to the "your request has been received" message. Not to closing the chat window. To the result.

The company and the consumer usually run different stopwatches.

The company's stopwatch starts when you send a message and stops when the bot replies. That comes to two seconds. A wonderful figure for a report.

Your stopwatch started earlier. First you looked for where to turn. Then you explained the problem to the bot. Then you went in circles through the menu. You waited for a staff member. You repeated the whole story to them. You checked whether the problem was really solved.

Two seconds passed on the company's clock. Forty minutes disappeared from your life.

An instant reply and a fast resolution are different things. A bot can reply at lightning speed, but not to the right question. It can instantly send you to a help section you already read. It can tell you in a second that it cannot help you at all.

So measure not the speed of the AI's reply, but the time from need to resolution.

And pay separate attention to repeats. If a day later you turn to it again with the same problem, the previous request cannot be counted as successful. The bot closed the dialog. You did not close the problem.

The third number: your costs.

This includes not only the price of the service itself.

Count the fees, paid calls, trips, printing of documents, extra subscriptions, and other costs that appeared because of the new way of serving you.

If the bot did not resolve the question and you had to travel to the office, add the cost of the trip. If you had to take time off work for the visit, add the lost time. If a wrong answer led to a missed deadline or an unnecessary purchase, count that too.

Nerves are harder to enter in a table. But they have a price too. Usually the company learns of it later, when we leave a review, cancel a purchase, or move to a competitor.

Now compare your three numbers before the AI and after.

Three outcomes are possible.

First: actions, time, and costs went down. That means the AI genuinely improved the service. Which model runs inside and how many parameters it has does not matter much to us. We got value. Thank you, we will keep using it.

Second: nothing changed for us, but the company cut its costs. That is an internal win for the company. It has a right to exist. But do not tell us the bot appeared purely for our convenience. We know how to count.

Third: the company started spending less, and we more. Before, a staff member resolved the question in five minutes. Now we struggle with the bot for fifteen minutes and still call a staff member. The company saved on service, and we paid for that saving with our time. Such AI cannot be counted as useful to the consumer.

Remember a simple rule:

AI improves a service only when it reduces our actions, time, or costs, without worsening the other numbers.

Do not compare the new bot with the perfect picture from the ad. Compare it with the way you used to get the service before.

And do not keep the result to yourself. If your numbers went up, tell the company. It may genuinely not know what its new service looks like from our side of the screen.

 Send this to the company

On [date] I used the service [name] through the AI bot [name or website address].

I wanted [desired outcome], but ran into a problem: [brief description]. I had to take [number] actions, spend [time], and [extra costs]. I had to repeat my question [number] times.

Before, getting this service took [time and number of actions]. After the AI was introduced it became [longer, harder, or more expensive] for me.

Please review the consumer journey, remove the obstacles listed, and tell me how I can get this service more simply or reach a staff member.

1.2. Consumer Needs. Demand Better.

Steps, time, and money can be counted. But they do not fully capture the quality of an AI service.

Using AI should be convenient for us. We should understand what is happening. We should be treated with respect. These are not the whims of a spoiled customer; they are our basic consumer needs.

Do not excuse bad service with the line that "it is only a bot." The company chose the bot. It built the bot into serving us. It is responsible for what happens on our side of the screen.

Our demands fall into three groups: functional, informational, and emotional.

Demand convenience

A convenient place. Corporate AI may live on a website, in an app, a messenger, or a phone line. It should be easy for us to find it, open it, and use it.

If, just to ask a simple question, you are made to install an app, register, or pass through ten menus, do not treat this as inevitable. Tell the company that the chosen channel is inconvenient for you. Demand an alternative way to get the service.

Convenient timing. A round-the-clock bot only makes sense when it can actually solve the problem. The message "contact us during business hours" is not round-the-clock service.

Count all the time from the moment a need arises to the result. If the bot answers in a second but runs you in circles for half an hour, write it exactly that way: "Your bot replied instantly, but resolving my issue took me 30 minutes."

A simple way to get the service. We are not obliged to do the company's work. Re-entering data we already provided, copying details from the chat into a form, hunting for another section of the site, and telling our story all over again to an employee.

Every extra step must be named. Do not just write "the bot is inconvenient." Write: "I had to repeat the order number four times, go through three sections, and then explain the problem to the operator again."

A link to a human. We have the right to demand a clear and accessible way to reach an employee. Especially when it concerns money, health, work, housing, insurance, or another significant matter.

Do not spend half an hour guessing the secret word "operator." Demand once to be connected to a human. If the bot refuses or keeps answering in circles, record this and contact the company by another means.

Demand information

Explaining what the AI can do. Before we start, we should be told what the bot can do, what it cannot do, and what actions it can take. Do not guess whether the bot processed a refund or merely logged a message. Ask directly: "Have you made a decision or passed my request to an employee? What happens next, and by when?"

Demand a concrete answer. The phrase "your request is being processed" explains nothing.

Explaining how data is used. Do not share sensitive information until you understand why it is needed and what will happen to it.

If the bot asks for passport, banking, medical, or other personal data, put direct questions to the company: what information is collected, why, where it is stored, to whom it is passed, and whether the conversation is used to train the AI.

A forty-page privacy policy is no substitute for a clear answer to a specific question.

A reliable answer. Do not accept the claim that the bot's mistake is somehow not the company's mistake. We are not dealing with an abstract neural network, but with an organization that put it in place to talk to us.

If the bot named a price, deadline, or condition, save the answer. For a significant decision, demand confirmation from the company or an employee. If the information turns out to be wrong, demand correction and the restoration of your interests.

Feedback. After a complaint or an error report, demand a case number, a response deadline, and information about the outcome.

We should not have to guess whether anyone read our message. "Thank you for your feedback" is not the handling of a complaint. We need to know what happened, what was fixed, and what to do next.

Demand respect

Respectful communication. We are not obliged to put up with foolishness, pressure, accusations, and misplaced jokes just because we are talking to a machine.

If the bot makes you feel guilty, ignores the seriousness of the situation, or answers inappropriately, stop the pointless conversation and record it. Demand that the company have this exchange reviewed by a human.

Accuracy. Do not fix the same mistake endlessly. If the bot confused your name, address, amount, order, or the circumstances of the case, demand that the data be corrected not only in the current message but in the company's system. An "AI hallucination" explains the cause of the error, but it is not grounds to shift its consequences onto the consumer.

A real result. A closed chat does not mean a solved problem. If after the conversation you did not get the service, confirmation of an action, or a clear next step, do not accept that the case is finished. Demand that it be continued and carried through to a result.

Remember three simple demands:

It should be convenient for me.

It should be clear to me.

I should be treated with respect.

Do not hesitate to say them out loud. The more often we make such demands, the sooner companies will understand: bad AI cannot be forced on the consumer in the guise of progress.

Send this to the company

Subject: Demand to remedy shortcomings in AI-based service

On [date] I used the service [name] through the AI bot [name].

During the interaction I had to [unnecessary steps]. I did not receive clear information about [the bot's capabilities, the decision, the data, or the next steps] and could not [reach an employee / correct the error / obtain a result]. As a result I spent [time and costs], and my problem remained unresolved.

I demand that my request be reviewed with an employee involved, that I be given a precise answer on the merits, that the mistake be corrected, and that I be told what changes will be made to the AI bot's operation.

Please register the request and send me its number and the review deadline.

1.3. Demand Plain, Human Language

Bad language also wastes our time. "Your request is under processing." "The data provided is undergoing verification." "To effect a refund it is necessary to provide the details."

What exactly is happening? Who is reviewing the request? When will the money be returned? What do I need to do right now?

Corporate AI can create a convincing impression of an answer while telling us nothing useful. It can write three polite paragraphs after which we know exactly as much as we did before the conversation. Do not mistake wordiness for information. Demand a clear answer.

Demand that it identify itself

We should know who we are talking to: a human or an artificial intelligence. The company must not force us to guess it from strange answers and repeated phrases. Information about the use of AI should appear at the start of the conversation, not hide in the terms of service.

Ask directly:

Are you an AI bot? Does an employee check your answers? Can you make decisions on behalf of the company?

It is especially important to understand the AI's role when it concerns a price, a refund, a contract, a loan, insurance, work, health, or another significant service.

Demand plain language

We are not obliged to translate corporate officialese into human speech. Not "to effect payment," but "to pay." Not "a refund will be carried out," but "we will return the money." Not "within the established timeframe," but "by 6 p.m. on Friday." A normal answer should meet a few simple demands.

The main thing first. If the answer is "yes" or "no," let the bot start with exactly that. Explanations can come after. We are not obliged to read three paragraphs to find the answer in the last line.

Concrete deadlines and amounts. Not "shortly," but "within three business days." Not "a fee may apply," but the specific amount of the fee and the conditions under which it applies.

Plain words. Internal "tickets," "escalations," "scoring decisions," and abbreviations are the company's language. We need an order, a complaint, a check, a decision, and a response deadline.

Short actions. The bot should say clearly what is required of us right now: "send the order number," "confirm the address," "you need to do nothing, we will reply tomorrow."

An answer to our question. Not a general overview of the topic, but an answer on the merits. If we ask when the money will be returned, an account of the refund rules is not an answer.

If the wording is unclear, do not hesitate to demand an explanation. It is not for us to adapt to the machine's language. The machine was put there to serve us.

Demand explanations for decisions

The phrase "the system made a decision" explains nothing. If the AI denied a service, changed a price, rejected an application, restricted access, or offered a particular option, ask:

What decision was made? On the basis of what data? Which conditions did I fail to meet? How do I correct the error or appeal the result?

They do not have to reveal the model's design, the program code, or a trade secret to us. But we need a clear explanation of what happened to us specifically. "That is what the algorithm decided" is not an explanation. The algorithm does not work on its own. The company chose it and uses it.

For significant decisions, demand confirmation in a form you can keep. Save the conversation, take a screenshot, or ask for the answer to be sent by email. We will return to the promises of corporate AI in the next chapter, when we demand the AI receipt.

Do not learn to guess the magic words

If the bot did not understand a normal question, do not turn the conversation into free testing of the system. You can phrase the question differently once. But if after two or three attempts the bot keeps answering off-topic, the problem is no longer your wording.

Do not run through magic words: "operator," "human," "call a manager," "connect me to support." Write a direct demand and record the answer. We are dealing with a service, not sitting an exam in how to talk to a machine.

Demand a human

AI must have limits. If it cannot solve the matter, we must be able to continue the conversation with an employee. The path to a human should be short and clear. Not hidden behind seven buttons, an endless menu, and requests to describe the problem yet again.

Demand that the employee be given the whole history of the conversation. We have already told the company our name, order number, and the circumstances of the case. Repeating it all from scratch just because two corporate systems cannot exchange information is not our duty.

Be especially insistent on human involvement if the AI's mistake could lead to loss of money, denial of a service, or other serious consequences. Not every question has to be handled by an employee. But the option to reach one should be real, not decorative.

Five phrases for talking to a bot

Answer "yes" or "no" first, then explain the reason.

Explain this in plain words, without the company's internal terms.

Name the exact amount, the deadline, and the action required of me.

Explain what decision was made, on the basis of what data, and how to appeal it.

Connect me to an employee and pass on the whole history of our conversation.

 Send this to the company

Subject: Request for a clear answer and access to a human agent

On [date] I contacted the AI bot [name] about [the issue].

The bot gave no clear answer, used [vague wording or internal jargon], did not explain [the decision, deadline, amount, or next steps], and offered no way to reach an employee. I had to repeat the question [number] times, but the problem remained unresolved.

I demand a clear answer on the merits, precise deadlines and next steps, an explanation of the decision made, and that my request be passed to an employee together with the full history of the exchange. Please also tell me how other consumers can move from the AI bot to a human in a similar situation.

Chapter 2. Demand the AI Receipt

In the first chapter we learned to judge corporate AI by its value to the consumer. Did it get faster, simpler, and cheaper for us? Did we receive a quality, safe, and respectful service in plain language?

Now the next question arises: how do we prove what actually happened? As long as the conversation with the bot goes well, nobody needs proof. But if the AI quoted the wrong price, lost our request, denied a service, or gave advice that cost us money, the situation changes. The company sees its logs and internal records. What do we have left? At best a screenshot. At worst, nothing.

And it is no accident that we begin with the receipt. Of the three documents we will discuss, the receipt is the easiest to obtain. A service passport or an AI policy is not something a company creates in a single day. In many cases, however, a company may already be able to provide a transcript from chat history or system records. That record is not always retained or available, so ask the company to confirm whether it exists and provide it. Until a full receipt exists, even a screenshot will do. This is our first and most accessible tool.

That is why our next demand is:

A significant interaction with corporate AI should end with the issuing of an AI receipt.

2.1. The Receipt as the Gateway to Justice

In 1879 a saloon owner from Dayton, Ohio, James Ritty, together with his brother John, patented one of the first cash registers. The reason was prosaic: Ritty suspected that his bartenders were pocketing part of the takings.

The machine was named the "Incorruptible Cashier". It recorded sales, displayed the amount, and gave a signal when a transaction was registered. One of the commercial versions added a paper tape with records. Ritty was not fighting for the rights of buyers. He wanted to control his own employees and keep his money.¹

But out of this idea grew something far greater.

Every deal must leave a trace.

For nearly a century and a half the modern consumer infrastructure has been built around that trace: the price tag, the label, the cash register receipt, the warranty, the contract, the complaints book. All of these tools answer simple questions:

What were we promised? How much did we pay? What did we get? Who is responsible for it? Where do we turn if something goes wrong?

A receipt does not guarantee justice by itself. A shop can still sell a bad product, and a company can still refuse a return. But the receipt opens the gateway to the procedure for protecting our interests. With it we can demand a refund, make a warranty claim, file a complaint, dispute the amount, and prove that the deal took place at all.

Without a receipt, a dispute quickly turns into a squabble:

- "I was quoted a different price."
- "Our employee could not have said that."

- "But he did."

- "Prove it."

The receipt replaces an argument about memories with a record of facts. That is why it protects both sides. The buyer cannot attribute to the seller something that never happened. The seller cannot pretend that the deal or the promise was just a dream.

It was precisely the ability to record a fact that let strangers and organizations trust one another. We do not have to know the shop owner personally. We receive a document that confirms: the interaction took place and left an evidentiary trace.

Now companies are placing AI agents between themselves and us.

These agents advise, sell, pick out products, quote prices, take orders, offer terms, collect data, and sometimes take part in decisions that have serious consequences for us.

But the economics of AI still work in many ways like the old saloon before the cash register arrived. We talked to a bot. It promised something. Quoted a price. Asked for data. Suggested a solution. Then the window closed, the history vanished, and we were left with nothing in our hands.

The company may have kept a technical log. But we cannot see it. We do not know what exactly is recorded there, how long it is kept, or whether the record matches the conversation we saw.

The result is a strange situation: companies are creating the smartest staff in history but forgetting to put a cash register next to it. We are asked to take AI at its word. Only artificial intelligence has no word of its own. It speaks on behalf of the company, and responsibility for its promises must not vanish when the chat window closes.

That is why we need an AI receipt.

An AI receipt is not an ordinary cash register receipt and does not replace one. It is confirmation of a specific interaction between a person and corporate artificial intelligence.

It should make it possible to establish:

what the person asked or instructed;

what the AI answered and did;

what result or decision the person received;

when the interaction took place;

on behalf of which company the AI acted;

how to dispute the result or reach an employee.

If an interaction can have consequences for us, it should leave a trace we can access.

Not every conversation needs a big document. If we asked a bot what time the shop closes, a separate receipt may be excessive.

But if the AI takes an order, calculates a price, collects personal data, gives a significant recommendation, changes terms, issues a refusal, or affects our money and rights, a receipt becomes necessary.

The law already knows a similar logic. For example, in the European Union, when many distance contracts are concluded, the company must provide the consumer with confirmation on a durable medium that allows the information to be stored and reproduced.² Under the Digital Services Act, platforms are

required to give users an explanation of the reasons when they restrict or remove their content. This explanation helps to understand and dispute the decision.³

We can demand a receipt right now, even though the law has not yet set a single duty to issue one. We vote with our money, and a company that listens to us will make a proper AI receipt its advantage sooner than it is required to.

Perhaps the first company will reply that it has no such function. The second will send an ordinary export of the conversation. The third will think it over and add a "Download history" button. And the fourth will be the first to create a proper AI receipt and make it its competitive advantage.

That is how a new practice comes into being.

Send this to the company

Subject: Please add the issuing of an AI receipt

On [Date] I used the AI bot [name] to obtain the service [name of service]. The interaction went well, but after the conversation ended I did not receive any confirming document.

Please provide me with a copy of the full interaction history and add the option to automatically receive an AI receipt after significant requests. It should retain the dialogue, the result obtained, the date and time, details of the AI service, and a way to reach an employee or appeal the decision.

Even though everything went well this time, such a document is necessary in case of an error, a dispute, or other consequences. Please let me know whether I can obtain an AI receipt for my request and whether the company plans to introduce this option.

¹ Patent No. 221,360 of James and John Ritty was granted on November 4, 1879. [Library of Congress](#), [patent text](#).

² Explanation of Article 8(7) of the EU Consumer Rights Directive: confirmation of a distance contract must be provided on a durable medium. [EUR-Lex](#).

³ Article 17 of the Digital Services Act requires hosting services to provide users with a clear and specific statement of the reasons for restricting content. [European Commission](#).

2.2. What an AI Receipt Should Contain

Let me clarify right away, to avoid misunderstanding. There is no single law today that would require every company to issue a full AI receipt after a conversation with a bot. But the law is already moving in that direction, and it is useful to know this.

Separate elements of this structure are appearing in legislation. In Quebec, Law 25 requires informing a person when a decision is based exclusively on the automated processing of personal data. On request, the person can find out what data was used, what reasons and main factors led to the decision, correct the data, submit observations, and reach an employee able to review the decision.¹

In Utah, requirements are in force to disclose the use of generative AI in certain consumer interactions. Consistent, clear disclosure can give a company limited protection against a specific enforcement action, but does not release it from responsibility altogether.²

In the European Union, logging requirements relate primarily to high-risk AI systems and depend on the organization's role. The GDPR lets a consumer obtain a copy of the personal data being processed but does not guarantee the automatic release of the entire technical log.³

These rules confirm the direction of movement: a person should know that a decision was made by AI, what data was used, and how to seek a review. But none of them yet creates an AI receipt in full.

So next we offer a consumer view, a draft AI receipt assembled from rights that already exist, the ordinary practice of issuing receipts, and common sense.

What the receipt contains

Who provided the service and to whom. The company name, the name and version of the AI assistant, the date, the start and end time of the interaction, and the details of the service recipient.

What service was provided. For example, a price calculation, placing an order, matching an offer, reviewing an application, or a consultation.

The full conversation. The entire dialogue in a storable text format, for example Markdown. Not a brief summary prepared by the AI, but the original dialogue in full.

Data obtained and consents. What data the AI requested and received, what we consented to and when.

The AI's actions. What exactly it did: the amounts quoted, the terms calculated, the deadlines set, and the next steps proposed. Itemized, like the lines on an ordinary cash register receipt.

The AI's authority level. What the AI was authorised to do in this conversation: only inform, recommend, carry out an authorised action, or make a decision. And how it ended: automatically, prepared by the AI but decided by an employee, or checked by a person. We need to know not only what the AI did, but whether it had the standing to do it.

The result of the service. What was arranged, calculated, recommended, changed, or rejected.

The price paid by the consumer. The price, fees, and any other charges shown to or paid by the consumer, if any. This does not mean the company's internal cost of providing the service.

Errors and employee intervention. Failures, repeated actions, corrections, handing the conversation over to a human, and decisions made by an employee.

The rules in force. The version of the AI service passport and of the company's AI policy that applied at the moment of interaction.

The retention period and challenge deadline. How long the company keeps the records of this interaction and by when we can still challenge it. Evidence only helps while those records exist.

How to file a complaint. Where to turn if we disagree with the result or find an error.

A direct line to the company. A visible button or link that lets us continue the conversation with an employee.

The receipt should include not just a link to the AI policy and the service passport, but a snapshot of the version that was in force during the conversation.

This is fundamental. A link leads to a page the company can change tomorrow. The receipt must preserve our agreements in the form in which they existed at the moment the service was received.

An AI receipt should exist in two identical copies. One stays with us, the other with the company. Both are generated from a single record of the interaction.

But two files are not enough if either of them can be quietly edited. So the receipt should have a unique number, a server timestamp, the versions of the policy and the passport, and a link or another way to verify its authenticity.

The facts in the receipt must come from authoritative system records: timestamps, the original conversation, actions, consents, results, and the versions of the policy and service passport. Generative AI may assemble these records and produce a plain-language summary, but it must not invent or reconstruct them from memory.

Because a receipt may contain financial, medical, identity, or other sensitive data, it must be delivered securely. It should apply data minimization, contain only what is needed as evidence, protect access to the file, and redact personal data belonging to third parties.

For significant interactions, a hash or electronic signature may be used. If a correction is made to the receipt, a new version is created, and the previous one remains in the history.

Now the dispute over "what the bot actually said" is settled simply: the parties compare their copies.

Such a receipt is our consumer standard: what we demand today and what tomorrow will become the norm of the market.

And until companies learn to issue such receipts, keep the evidence yourself. Download the conversation. Take screenshots of the key messages. Record the date, time, bot name, and page address. Save the amounts quoted, deadlines, promises, refusals, and attempts to reach a human.

If the interaction had serious consequences, immediately send the company a letter describing what happened and ask it to confirm the AI's answer or action. Do not wait for the chat history to disappear.

Today we take screenshots. Tomorrow we should receive a proper AI receipt.

 Send this if there is no AI receipt

Subject: Please provide a record of the interaction with the AI

On [Date and time] I used the AI bot [name] to obtain the service [name]. As a result of the conversation the bot [made a decision, quoted a price, arranged an action, or gave a recommendation], but after the interaction ended I did not receive an AI receipt or a full record of the conversation.

Please provide the full history of the dialogue, details of the actions performed and the result obtained, the data used and the consents given, the service rules in force, and the way to appeal or reach an employee.

If the company does not yet generate AI receipts, please retain the logs of this interaction, confirm the result described, and let me know whether the automatic issuing of such documents to consumers is planned.

¹ [Quebec Act respecting the protection of personal information, Article 12.1.](#)

² [Utah Code, Title 13, Chapter 77.](#)

³ [EU AI Act, Articles 12, 19 and 26\(6\); GDPR, Article 15\(3\).](#)

Chapter 3. Demand the AI Service Passport

The AI receipt records the practice. It shows what we asked, what the bot answered, what data it collected, what actions it took, and how the interaction ended.

But the receipt alone still does not answer the main question: what should the service be like?

Say the bot solved the problem in twenty minutes. Is that fast or slow? It asked for a photo of a passport. Did it have the right to request it? It refused a refund. Was such a ground provided for? A human joined after the fifth request. And when should they have joined?

To judge a fact, we need a norm. To check the real service, we need to know in advance how it should be provided. That norm becomes the passport, or the AI service standard.

The passport says what we were promised. The AI receipt shows what we got.

3.1. The Passport as a Yardstick of Quality

Service standards are nothing new.

We are used to a product having a description, a delivery having a deadline, a tariff having terms, a warranty having rules, a public service having a procedure. We want to know in advance what we will get, how long it will take, how much it will cost, and what to do if a problem arises.

Corporate AI should not become an exception.

If a company uses a bot to provide us with a service, that service must also have a standard. Not an internal instruction for the developer, but a public promise to the consumer.

The passport does not tell us which model runs inside, how the system prompt is written, or which databases the bot is connected to. Those are the company's internal workings. What matters to us is different:

What will we get?

By what deadline?

What can the bot do?

What is it not entitled to do?

What data will it request?

When will it call a human?

Where do we complain?

Without a passport, a company can always say the bot "worked as intended." Only we do not know how it was configured or what we were promised at all.

Imagine the AI receipt reports: a human joined after eighteen minutes. Is that a lot or a little?

If the passport promised a human within two minutes, we have a breach of the standard. If no deadline is set anywhere, all that is left is to argue about impressions.

Or the bot requested medical or banking data. The AI service passport must state in advance an exhaustive list of the data needed for the service. Then we can check whether the bot had the right to ask such a question.

Or the AI refused the service. The AI service passport must contain a closed list of grounds for refusal. Then we compare the reason stated in the receipt with the published list. Is the reason for refusal in the passport, or did the bot invent it in the course of the conversation?

So the subjective "it seems to me the bot worked badly" turns into a verifiable statement:

The passport promised one thing. The AI receipt recorded another.

That is precisely why the passport must exist before our interaction, not appear after a complaint. A company must not change the rules retroactively to explain a mistake already made.

The AI service passport must have a version and a date. Updated the bot, changed the price, added a new feature or removed an old one: release a new version. The previous one must remain in the archive.

The AI receipt records exactly the version of the passport that was in force during the conversation. If the interaction took place in March, it must be judged by the March standard, even if by April the company has already fixed everything.

There is a simple set of three documents:

The policy explains how the company uses AI in general. The passport sets the standard for a specific service. The receipt records a specific interaction.

One does not work without the other. A receipt has nothing to compare against without a passport. A passport cannot be checked without a receipt. And without an AI policy it is unclear who answers for the whole system.

The law does not currently require every company to publish a passport in exactly this form. Specific legal duties depend on the country, the industry, the data, and the nature of the service.

But service standards are normal consumer practice. We have the right to ask a company to state in advance what exactly it promises us. And if a significant service is provided by AI, such a promise is especially important.

Do not settle for the phrase "our bot helps customers." Ask: how exactly does it help, what does it guarantee, and where do its powers end.

3.2. What the Passport Should Contain

The AI service passport is not for beauty and not for reporting. It must answer the consumer's practical questions.

Here is what we want to see in the passport of every significant AI service.

What this service is. The name, the purpose, the company that provides it, the way to obtain it, and the range of people to whom it is available.

What result we will get. Not "the bot will help solve the issue," but a concrete result: it will calculate the cost, select an offer, process a refund, accept an application, or book an appointment.

By what deadline the service will be provided. The bot's response time, the overall time to obtain the result, and the time for a human to join.

What the bot can and cannot do. A full list of actions and honest boundaries. For example: the bot can process a refund for a standard product, but does not handle payment disputes.

Boundaries are especially important. We should not spend twenty minutes on a conversation only to learn at the end that the bot had no right to resolve our issue from the very start.

What data we will be asked for. An exhaustive list of the required information. Anything not on this list requires a separate explanation.

How much the service costs. The price, fees, and other mandatory charges. If the service is free, that too must be stated directly.

On what grounds we can be refused. A closed and clear list of reasons. Not just "the application does not meet the requirements," but which requirement exactly was not met and what can be fixed.

When a human is obliged to join. For example, if we asked for it directly, the bot failed twice to understand the question, the case is not covered by the standard, an error was found, or the conversation turned into a conflict.

The handover to a human must not depend on the algorithm's mood. This is our emergency exit, and we want to know where it is.

What is recorded during the interaction. Which actions, decisions, data, consents, and errors are logged. What AI receipt we will get after the service is completed.

How the result is checked. What happens if the bot is unsure of the answer or found a contradiction. In which cases the result must be checked by a human before we receive it.

Where to complain. The channel of contact, the time for review, the way to get an answer, and the procedure for a human to reconsider the decision.

This is our consumer demand for the passport, not a mandatory form from the law. We ask companies to answer these questions in advance and publicly, because it is from such demands that a new service practice takes shape.

If consumers consistently ask about deadlines, the AI's boundaries, data, responsibility, and the appeal procedure, companies will start fixing the answers in standards. A repeated practice will become a market norm, and over time may find reflection in industry standards and legislation.

The passport must be public. It should be placed next to the bot or on the service page, not inside a forty-page user agreement.

It must be written in plain language. Not "the request will be reviewed in the established manner," but "we will answer within three working days."

And most important: what is written in the passport must match how the AI actually works.

If the company promised to connect a human within two minutes, we have the right to time it. If it stated that the bot does not request banking data, it must not request them. If it named a closed list of reasons for refusal, a new reason cannot appear out of thin air.

The bot represents the company. Its price is the company's price. Its promise is the company's promise. An error of the algorithm does not cancel responsibility to the consumer.

The passport must be available not only to people but also to personal AI assistants. After all, what do we really want? Not to run around ten sites ourselves comparing terms, but to say: "find me the best and prepare it." We already do this: we ask ChatGPT to select a service, compare offers, check terms. There

will only be more of this going forward: our agents will start reaching out to companies' agents themselves, requesting information and preparing deals. And that means people and companies will have to agree in advance on the rules by which these agents communicate. That is exactly what clear rules, service passports, and honest receipts at the end are for: without them our agent simply will not understand who it is dealing with, and will go to whoever described itself more clearly.

For such interaction, a pretty page on a website is not enough. Our agent must quickly determine what services the company provides, what its AI can do, what data will be required, how much the service costs, what the deadlines and limits are, whether an AI receipt is issued, and where to send a request.

If this information is hidden in advertising copy, PDF files, and user agreements, a personal agent finds it hard to locate and match. The interaction may simply not happen: the agent will not understand the terms and will choose the company whose services are described more clearly.

So demand that the passport be published simultaneously in two forms: an understandable document for a human and a structured machine-readable format for AI. The content must match, and the version number and update date must let you check that it is current.

This benefits not only consumers. A company with machine-readable passports becomes visible to a new class of customers who will come through personal AI agents. A personal agent is likely to recommend first the company it can understand and transact with quickly. It will probably skip a company whose services it cannot interpret.

Demand the passport before a significant interaction begins. Keep its version together with the AI receipt. And if there is no passport, put to the company the questions this document should have answered.

 Send this to the company

Subject: Please provide the AI service passport

I used the AI service [service or bot name], but could not find its public passport or standard of provision.

Please tell me what result the service guarantees, within what deadlines it must be provided, what the AI bot can and is not entitled to do, what data it may request, how much the service costs, on what grounds a refusal is possible, and in which cases a request is handed over to a human.

Please also state what is recorded during the interaction, what AI receipt the consumer receives, and how the result can be appealed.

If no such document exists yet, please develop and publish an AI service passport next to the bot, stating the version number and the date it takes effect. This will let consumers understand the service terms in advance and compare the promised standard with the actual result.

Chapter 4. Demand the AI Policy

We now have two documents. The passport tells us how a specific AI service is supposed to be provided. The receipt records what happened during a specific interaction.

But a company may have more than one bot and more than one service. AI advises customers, recommends products, takes requests, calculates prices, checks documents, and answers complaints. Each service needs its own passport, but some rules will repeat.

Who is responsible for the AI? How does the company handle our data? How do we reach a human? Where do we complain? What happens when there is an error?

The answers to these questions should not differ from bot to bot. They need one shared document: the company's AI policy.

4.1. Why a Company Needs an AI Policy

Before we use a corporate AI, we want to know how everything works here in the first place.

Not in one specific service, but across the whole company. Which AI systems interact with us? How do we know we are talking to a machine? Who is responsible for how they work? What happens to our data? Where do we find a human?

The answers should be gathered in one public document on the company's website. It might be called "Artificial Intelligence Use Policy," "Our AI Rules," or something else.

The name does not matter. What matters is that the document is easy to find and understand.

This is not a replacement for the privacy policy, the terms of service, and other legal documents. It is a clear entry point into them. Two or three pages in plain language, answering the main questions and linking to the detailed documents.

Part of this transparency the law already requires. I spoke about this in the introduction: [Article 50 of the EU AI Act](#), effective 2 August 2026, requires warning a person that they are talking to an AI when it is not obvious, and adds separate duties for certain systems and content.¹

The handling of personal data is governed by the GDPR regardless of whether the company uses artificial intelligence or not. The company must define the purposes and legal grounds for processing, the categories of data, the retention periods, the recipients of the information, and the person's rights. The specific duties depend on the situation and the company's role.²

The EU AI Act and the GDPR do not require gathering all of this into one document in exactly our form. The AI policy becomes the frame into which the company puts both the requirements of the law and the voluntary guarantees to the consumer.

This benefits more than just us.

Imagine a small company with two AI bots. The support contacts can be written by hand into each bot and repeated in two passports. For now, everything works.

A year later the company already has twenty AI services. The email address has changed. A new responsible person has been appointed. The complaint review times have been updated. New data-handling rules have appeared.

Now the information must be changed in twenty passports, twenty bots, and dozens of website pages. Somewhere an old address, a wrong deadline, or the name of an employee who left long ago is bound to remain.

That is why a simple principle is used: a single source of truth.

General information is kept in one place. The company updates it once, and the service passports and AI bots use the current version. A snapshot of the policy in force at the time of the conversation is included in the AI receipt.

When a support contact changes, the company corrects one entry. After that, the new contact appears in all linked passports, bot answers, and receipts.

For us this means that the company's different AIs will not give different answers to the same question. For the company it means fewer errors, disputes, and manual work.

The AI policy becomes at once a promise to the consumer and a single reference book for the whole system.

The company has as much interest in this transparency as we do. A clear AI policy strengthens the trust of today's consumers and makes the company accessible to tomorrow's agent economy.

4.2. What the AI Policy Should Contain

The AI policy should answer the consumer's simple questions.

Who is responsible for the use of AI. The name of the legal entity, the company's contact, and the responsible unit or employee.

We do not need the internal management structure. But we should understand who is able to answer a question, review a complaint, and correct an error.

Where the company uses AI. A list of the AI services and bots that interact with consumers.

Next to each service there should be a link to its passport. Then we first see the overall picture, and then move to the terms of the specific service.

How the company discloses the use of AI. When and in what form the bot introduces itself, how AI interaction is labeled, and how to tell an automatic answer from an employee's answer.

How our data is handled. What information is collected, for what purpose, on what grounds, to whom it is passed, and how long it is kept. Are conversations used to train or improve models? How to get a copy of your data, correct it, or delete it, where such a right applies?

The detailed legal terms may stay in the privacy policy. But the AI policy should give a clear explanation and a direct link to the relevant section.

Whether the AI makes decisions about people. What decisions may be made or prepared automatically? How to learn the reasons for a specific outcome? How to correct the data used and demand human review, where this is provided by law or by the company's rules?

How to reach an employee. Contacts, the procedure for handing over the conversation, and the promised response time.

The ability to demand a human is not a universal legal right for every conversation with every bot. But we suggest that companies adopt it as a voluntary guarantee, especially for significant services and disputed situations.

What happens when there is an error. How the company corrects a wrong answer, reverses an erroneous action, restores the consumer's interests, and reports the outcome.

What traces remain. What is recorded during the interaction, how long the information is kept, and what AI receipt the consumer gets.

Where to complain. The channel for filing a request, the review period, the way to receive an answer, and the appeal procedure.

How the policy changes. The version number, the effective date, and a short description of the changes.

Previous versions should stay available. An interaction that happened in March is assessed against the March policy, not against a document the company rewrote after a complaint in April.

A register of trust credentials can be added to the AI policy: licenses, certificates, adopted standards, quality marks, results of independent audits, and public test reports.

This is an additional section, not a mandatory part of the universal form. It lets the company not just claim reliability, but show what backs it up.

Here it is worth clearly separating what is already mandatory by law from what for now remains our voluntary standard.

Informing about interaction with AI, transparency of individual automated decisions, and rules for handling personal data in applicable cases may already be mandatory. A single list of all AI services, a universal path to an employee, the issuing of AI receipts, a public register of credentials, and the very form of the shared document remain, to a large extent, a voluntary standard for now.

But voluntary does not mean unnecessary.

If we consistently ask companies about their AI, data, responsibility, and error-correction procedures, the answers will start to gather into an AI policy. Repeated practice will become a market norm, and then it may be reflected in industry standards and legislation.

A specific service's passport may refine the AI policy, but must not quietly weaken the guarantees. If the company promised a fast path to an employee in the AI policy, a separate bot must not turn that path into an hour-long wait.

And the whole construction must work as a single whole:

The policy says how the company works. The passport says what is promised. The receipt records what happened.

Send this to the company

Subject: Request to disclose your artificial intelligence use policy

I use your company's services and see that you employ AI bots and other AI systems to work with consumers. This can make service more convenient, but I could not find any single AI policy for their use.

Please tell me which AI services interact with consumers, who is responsible for how they work, what data they collect, and whether conversations are used to train models. Please also explain whether the AI makes decisions about customers, how to reach an employee, get an AI receipt, correct an error, or file a complaint.

I suggest publishing this information in a single document on the company's website, stating the version number and the update date, and also providing a machine-readable version for consumers' personal AI assistants.

¹ [EU AI Act, Article 50](#); [European Commission guidance](#).

² [General Data Protection Regulation](#).

Bonus: a voluntary transparency label

One last tool for this chapter. Here are eight ready-made "AI works here" stickers that you can pass to a company. Ask it to display a clear notice wherever you meet its AI: on the website, in the chat widget, on the office door, or in the newsletter footer. The point is simple: you should be told honestly and visibly that you are interacting with a machine rather than a human.

These stickers are voluntary transparency labels. They are not official EU marks and do not by themselves prove legal compliance. They are simply examples of what a clear AI notice can look like. Consumers can send them to companies with a suggestion: use one, adapt it, or create your own, but make the use of AI visible and understandable.





✉ **Send this to the company**

Subject: A request to disclose AI use honestly

I use your services and I can see that you use AI to interact with consumers. Please disclose this honestly and visibly: place an "AI works here" notice wherever I meet your bot, on the website, in the chat widget, in messengers, in the app, and at the entrance to your office.

Eight examples are included in this brochure. They are voluntary transparency labels, not official marks. You may use one, adapt it, or create your own. Please let me know where and when you will display a clear AI notice.

Chapter 5. Change the System to Fit You

We have worked out what we have the right to demand from corporate AI. A proper service should be useful, understandable and convenient. A meaningful interaction should end with an AI receipt. Every service should have a passport, and every company should have an AI policy.

But these documents will not appear on their own. The rules of corporate AI are taking shape right now, out of millions of daily conversations between people and bots. Every request, demand, suggestion, and even a silent departure shapes what these services will become.

That is why this final chapter is about action.

5.1. First, Stand Up for Yourself

If the AI made a mistake, do not start with reflections on the future of humanity. First, solve your own problem.

Demand that the data be corrected. That the decision be reviewed. That the money be returned. That the promise be kept. That a human employee be brought in. That the reason for a refusal be explained. That an AI receipt be provided.

In the previous chapters we gave ready-made templates for complaints. Use them. Insert the name of the company, the service, the date and the circumstances of your case. Attach the correspondence and screenshots.

Now we have three tools that make a complaint far stronger.

The AI receipt records what happened.

The passport shows what the company promised.

The AI policy explains who is responsible and how mistakes are corrected.

That is why a good complaint can be very short:

The service passport promised [X]. The AI receipt records [Y]. The AI policy says [Z]. Please correct the result, explain why the difference occurred, and tell me what will change so it does not happen again.

Describe exactly what happened:

what task you wanted to solve;

at which step the problem arose;

how much time and effort you lost;

what the AI answered or did;

what result you are demanding;

what should change so the problem does not happen again.

Criticize the right way. Complain the right way. Demand corrections the right way.

We are not obliged to work for free as testers of corporate systems. But if a company has put an AI between itself and the consumer, it must know how it works on our side of the screen.

We are the first to notice that a bot does not understand a plain question, names the wrong price, requests unnecessary data, or hides the human behind seven buttons.

Do not assume your complaint will change nothing. One complaint may stay a private case. Ten similar complaints become a problem. A hundred turn into a demand from the market.

Our first and strongest lever is the market: ask for a correction, leave a factual review and, if necessary, take your money and loyalty to a competitor. But do not stop there when the consequences are serious or the company refuses to act. Escalate step by step: the company's complaints service; a consumer protection body; a data protection or sector regulator where relevant; an alternative dispute resolution body or a court. Use the channel that fits your case, keep the AI receipt and correspondence, and describe the facts calmly and precisely.

So ask not only to fix your case, but also to remove the cause.

If the bot lost your data, restoring your request is not enough. The way it works must change so it does not lose other people's data.

If you had to ask for an employee ten times, finally connecting you to a human is not enough. The path to a human must be shortened for everyone.

If the AI gave a wrong answer, sending you an apology is not enough. The instruction, the service passport, or the very rules of the bot's operation must be corrected.

Remember:

An apology fixes one case. Changing a rule removes the cause.

A company that spends a second year apologizing for the same mistake is fixing nothing.

So after the reply, ask: what exactly changed? Has the bot been updated? Has the service passport been corrected? Has the AI policy been supplemented? How will the company protect other consumers from the same problem?

Demand. Demand. Demand.

This is how we protect our own interests and, at the same time, shape a new standard of service.

5.2. From Feedback to Co-authorship

A demand and a complaint are the first step of participation. But not the last.

We can tell a company what was done badly. Or we can go further and explain how to do it better.

If you liked part of the service, say so. If you know a more convenient way, suggest it. If the company invites you to an interview, a test, or a focus group, take part.

This work has its own ladder.

First step: tell them about your experience

This is the simplest level.

Answer a short survey after the interaction. Leave a substantive review. Write whether the problem was solved, how long it took, and whether it became easier compared with the old way.

Do not just leave a star. Explain the reason for your rating.

For example:

The bot quickly found my order, which is convenient. But it could not change the address and sent me to the same section three times. Add a direct handover to an employee that keeps the conversation.

Such a review shows at once what should be kept and what should be fixed.

For simple services this may be enough. If a bot reports opening hours or a delivery status, a company does not need a standing consumer council. It needs to read feedback regularly and fix the obvious mistakes.

Second step: help test

The higher the consequences of a mistake, the more important consumer participation is before a new version launches.

If the company invites you to test a prototype, agree. Take part in interviews and focus groups. Join beta-tester clubs. Check new features before everyone is forced to use them.

The consumer has an advantage over an internal tester: we do not know how everything "is supposed to work".

We start a conversation in the middle. We write with mistakes. We ask two questions at once. We press the wrong button. We look for a service somewhere other than where the developer expected us to look. That is exactly why we find problems the company's team does not notice.

Five live consumers before launch can prove more useful than a hundred pages of an internal report.

It is especially important to test AI that sits on the customer's main path: taking orders, handling returns, working in the personal account, reviewing complaints, or taking part in meaningful decisions.

Third step: become a co-author

For AI that enters our life for the long term and affects money, work, housing, health, or access to services, a one-off survey is not enough.

Here you need a permanent connection.

It could be a consumer panel that is first to see new versions of the bot. A beta club with early access. Regular meetings with the company's team. A public or customer council. A joint review of the service passport and the AI policy.

The name does not matter. What matters is that consumers take part in editing the rules by which they are served.

Especially needed are people who can represent the experience of a specific group: the elderly, people with disabilities, migrants, parents, patients, borrowers, or workers. What is convenient for the majority can create an insurmountable barrier for others.

If you know such experience from the inside, do not wait for an invitation. Offer the company your help.

The AI policy cannot be written once and for all. Technologies change. Laws are updated. Our expectations grow. A new model can make an old bot behave differently.

That is why documents can be living or dead.

Living rules have a date of update. They are reviewed after a change in technology, a surge of complaints, the launch of a new service, or the arrival of a new law. Consumers take part in the review.

Dead rules were written once, posted on a website, and forgotten. The bot keeps changing, while the document still describes something long gone. Help keep the rules living. Report discrepancies. Suggest changes. Ask when the passport was last reviewed. Demand that real cases be taken into account.

Right now there is a window of opportunity for this. In a few years the rules of corporate AI will settle, turn into familiar procedures, and grow expensive infrastructure. Changing them will be far harder.

Today companies are still looking for solutions. And we can take part in creating them.

 Send this to the company

Subject: A suggestion to improve your AI service

I used the AI service [name] and want to share some feedback.

I liked [what works well]. At the same time, I ran into [the specific problem]. I suggest [how the service could be improved], because it would let consumers [what people would gain].

I am ready to tell you more about my experience, take part in an interview, a focus group, or a test of a new version. If the company is creating a consumer panel, a beta club, or a working group to review AI passports and rules, please tell me how to join it.

Please pass my suggestion to the employee responsible for developing AI services, and let me know whether it will be considered.

Conclusion. The Rules Depend on Us

And that is all. We have gone from a simple conversation with a corporate bot to taking part in creating the rules by which it will serve us.

Let us look at what we now have.

We have learned to check whether AI truly makes a service better. Not by the company's advertising and not by the bot's response speed, but by our own measure: actions, time and costs.

We have brought back into digital service one of the basic consumer practices: always take the receipt. It does not guarantee that the company will admit the mistake at once. But it confirms that the interaction took place, what we were promised, and what actually happened.

The receipt records the practice. But a practice can only be judged against a standard. That standard becomes the passport, or the AI service standard. Without a standard published in advance, a company can always say the bot "worked exactly as it was set up". The passport turns vague words into a checkable promise.

When AI services become many, repeating the same general information in every passport is pointless and dangerous. That is why a company needs a unified AI policy. In it we should find a list of AI services, the people responsible, the rules for handling data, the order of notification, the handover to an employee, the correction of mistakes, the issuing of receipts, and the review of complaints.

First, protect yourself. Demand that the mistake be fixed, the decision reviewed, the money returned, the information provided, and a human employee brought in.

Then demand that the cause be removed. An apology fixes one case. Changing a rule protects everyone who comes after us.

Criticize. Leave reviews. File complaints. But if you know how to do it better, suggest it.

Take part in surveys, interviews, focus groups, and tests of new versions. Join customer panels and beta clubs. Represent the experience of people whose needs the company might have missed.

Become co-authors of the rules by which you will be served.

Remember our structure like the multiplication table:

The policy says how the company works. The passport shows what is promised. The receipt records what happened.

One without the other does not work.

Without an AI policy it is unclear who is responsible for the AI, how our data is used, and where to turn when there is a problem.

Without a passport we have no standard to compare the result against.

Without a receipt it is impossible to prove what happened in a particular conversation.

Together these documents turn corporate AI from a black box into a proper consumer service.

Part of this structure is already backed by law. But usually consumer practice appears before the law. First people start to demand. Then companies that listen to their customers turn the demand into a

competitive advantage. After that the practice becomes a market standard, enters industry norms, and gradually takes legal form.

The rules are being written right now.

In a year, two, or three, the practice will settle. Familiar interfaces, standard contracts, industry solutions and new legal norms will appear. What companies create now will be reproduced for many years to come.

Changing an established system is always harder than taking part in its creation.

So do not hand this historic opportunity over to developers, technology vendors, lawyers and chief financial officers alone. They have their own interests. We should have ours.

There are billions of us. We pay for services with money, time, attention, data and loyalty. We have the right to demand that corporate AI be set up not only to save the company money, but also to meet our needs.

Demand the receipt.

Demand the passport.

Demand the policy.

Defend your rights.

If we stay silent, the rules will be written without us. If we act together, they will start to take our interests into account too.

Thank you for reading this brochure to the end.

I would be glad to hear about your experience with corporate AI, and to receive comments and suggestions. Tell me what is missing here, which demands worked, and how the companies responded.

My other brochures and materials can be found at aibusiness.vc.

For feedback, suggestions and collaboration: info@aibusiness.vc

About the Author

Sergei Ponomarev holds a PhD in political science. For more than 20 years he has been making services open, understandable and convenient for the people who use them: first in the public sector, now in business. For 7 years he ran a monitoring program for government and municipal services, conducted hundreds of independent quality assessments and mystery shopping audits, developed an Information Openness Standard for public authorities, completed a fellowship at the Kennan Institute in Washington, DC, and taught at a university for 14 years. For the past year and a half he has been building AI agents hands-on and running the portal aibusiness.vc.

Website: aibusiness.vc

Email: info@aibusiness.vc



Point your camera to open aibusiness.vc